

SCIENCE APPLIANCE

TAMING THE INCREASING COMPLEXITY OF LARGE-SCALE CLUSTER COMPUTERS

Computer clusters for running parallel tasks are traditionally assembled from off-the-shelf computers (called nodes), each computer with its own operating system. Starting, monitoring, and controlling a parallel task is time consuming and complicated, usually requiring separate attention to each node. The complex configuration has many possibilities for failure, needs extra software to solve problems, and is difficult to maintain.

The Science Appliance is a complete redesign of the traditional cluster architecture from the nodes up. It is a software suite, compiled by the Advanced Computing Laboratory, which includes LinuxBIOS, a basic input/output system that configures a computer from a cold start without a hard disk or other such device, a task many people in the computing community claimed was impossible.


The suite also includes the Beowulf Distributed Process Space (BProc), a set of modifications to the Linux Operating System, which provides a single-system image across all the nodes of a cluster. A single-system image makes the cluster appear as a single large computer and provides only one point of control—the front-end node—simplifying the management of parallel tasks. The user launches, manages, and views an entire job from the front-end node. This machine sends processes to the other nodes, which do not require their own software.

Combining BProc with LinuxBIOS results in an extremely reliable, purpose-built cluster node that reduces maintenance costs for the entire cluster to those for a single machine. Used today on some of the most powerful computers in the world, the Science Appliance is cutting-edge cluster architecture.




The first science appliance

The first Science Appliance was built for SC2000, the annual conference on high-performance networking and computing. It was composed of 13 slave nodes and one master node with one hard disk.



Platform for computer science research

Pink, a 1,024-node (2,048-processor) Intel Xeon Science Appliance, is an institutional computing platform designated solely for computer science research for its first 18 months at Los Alamos. Then it will be used for science research support.



Platform for secure computing

ASCI Lightning, a 1,408-node (2,016-processor) 64-bit Advanced Micro Devices Opteron cluster, is the largest Science Appliance built to date. Lightning is designated as a platform for secure computing.